

IZW
AF

- 1 -

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of:	:	Before the Examiner:
Edward A. Hubbard	:	Nilesh R Shah
Serial No.: 09/602,803	:	Group Art Unit: 2127
Filed: March 30, 2000	:	
Title: DATA SHARING AND FILE DISTRIBUTION METHOD AND ASSOCIATED DISTRIBUTED PROCESSING SYSTEM	:	United Devices, Inc. 12675 Research, Bldg A Austin, Texas 78759
Mail Stop Appeal Brief-Patents Commissioner for Patents P.O. Box 1450 Alexandria, VA 22313-1450		

APPEAL BRIEF**I. REAL PARTY-IN-INTEREST**

The real party in interest is United Devices, Inc. who is the assignee of the entire right and interest in the present Application.

CERTIFICATION UNDER 37 C.F.R. § 1.8

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to Mail Stop Appeal Brief-Patents, Commissioner for Patents, Alexandria, VA 22313-1450, on January 10, 2005.

Signature

Serena Beller

(Printed name of person certifying)

01/14/2005 SSE/SHE1 00000046 09602803
01 FC:2402 250.00 0P

II. RELATED APPEALS AND INTERFERENCES

There are no appeals or interferences known to Appellants, the Appellants' legal representative, or assignee which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

III. STATUS OF CLAIMS

Claims 23-30 are pending in the Application. Claims 23-30 stand rejected.

Claims 26 and 29 were amended for informalities. Claims 23-30 are appealed.

IV. STATUS OF AMENDMENTS

After Final amendments to Claims 26 and 29 were filed August 10, 2004 to put the claims in form for allowance but the Applicant has not received an Advisory action from the Examiner indicating the status of these amendments. The status of the after Final amendment is unknown.

V. SUMMARY OF THE INVENTION

In one embodiment, a method of operating a distributed processing system (FIG. 7A, element 100) having a network (FIG. 7A, element 102) coupling a multiplicity of Host distributed devices (FIG. 7A, elements 108, 110, 112) for processing workloads for the distributed processing system, a plurality of Client systems requesting processing of the workloads (FIG. 1B, element 130), and a Server system for selectively distributing the workloads from the plurality of Client systems for processing by the distributed processing system comprising five steps. Specification, page 5, line 13 through page 7, line 10.

In step 1, the Server system (FIG. 14B, element 104) receives a request from one of the plurality of Client systems (FIG. 14B, elements 108A-108H) to use the

distributed processing system to process a first workload. Specification, page 64, lines 13-16. FIG. 14A and FIG. 14B.

In step 2, the first workload is sent to a first Host distributed device selected from the multiplicity of Host distributed devices. Specification, page 68, lines 1-8; page 6, lines 6-13; page 6, lines 6-13.

In step 3, an index of one or more data addresses defining a location of first data required to process the first workload is sent to the first Host distributed device. Specification, page 67, lines 16-22.

In step 4, the first data from a first data address selected from the one or more data addresses in the index is accessed. Specification, page 66, lines 17-23.

In step 5, the index is updated to include a storage address of storage coupled to the first Host distributed device as a location of the first data. Specification, page 64, lines 13-23.

In another embodiment, a distributed data processing system comprises:

1) a multiplicity of Host distributed devices coupled to a network such that the Host distributed devices process workloads for the distributed processing system. FIG. 14B system 1450, elements 108A-108H.

2) a Server system (FIG. 14B, element 104) coupled to the network for distributing workloads to selected of the multiplicity of Host distributed devices. FIG. 8, element 102 and elements 108-112. Specification, page 5, lines 13-16.

3) a database coupled to the Server system for storing capability vectors having capability values defining an ability of each of the multiplicity of Host processing devices has for processing workloads for the distributed processing system. FIG. 6A, element 620. Specification, page 37, lines 1-23.

4) an index stored in the database having one or more storage addresses defining storage locations for accessing data required to process workloads for the distributed processing system. Specification, page 48, lines 8-15.

5) circuitry for accessing first data required for a first workload by a first Host distributed device processing the first workload using an address of the first data stored in the index, wherein the first Host distributed device stores the first data at a first data address when processing the first workload. Specification, page 48, lines 8-15.

6) circuitry in the first Host distributed device for automatically updating the index in the database to include the first data address as a location for the first data. Specification, page 66, lines 12-23.

In another embodiment, a computer program product operates within a Server system managing a distributed processing system, wherein the Server system is coupled to a network, the network configured to enable the Server system to be coupled to a multiplicity of Host distributed devices for processing workloads for the distributed processing system, the program product comprising a program of instructions for performing the program steps of:

1) configuring a database in storage coupled to the Server system for storing and accessing capability vectors having capability values defining an ability each of the multiplicity of Host distributed devices has for processing workloads for the distributed processing system. FIG. 6A, Specification, page 36, line 19 to page 37, line 4.

2) configuring an index in the database for storing addresses defining locations of data required to process each workload the Server system submits to the distributed processing system for processing. Specification, page 36, line 18 through page 37, line 17.

3) sending storage addresses of first data required to process a first workload from the index to a first Host distributed device when the first Host distributed device is selected to process the first workload. Specification, page 66, lines 14-23.

4) updating the index with a storage address of the first data within storage coupled to the first Host distributed device when the first Host distributed device is selected to process the first workload. Specification, page 66, lines 14-23.

In another embodiment, a software agent operates within each of a multiplicity of Host distributed devices coupled to a network, the network configured to enable a Server system to selectively couple the multiplicity of Host distributed devices to perform workloads for a distributed processing system, the software agent comprising a program of instructions for performing the program steps of (Specification, page 5, line 18 through page 6, line 4):

1) receiving in a first Host distributed device selected from the multiplicity of Host distributed devices a first workload, an application program to process the first workload, and an index of storage addresses defining one or more locations for accessing first data required to process the first workload. Specification, page 67, lines 16-22.

2) accessing the first data from one of the storage addresses in the index. Specification, page 66, lines 17-23.

3) storing the first data at a first storage address in storage coupled to the first Host distributed device when the first Host distributed device is selected to process the first workload. Specification, page 6, lines 15-23,

4) updating the index by adding the first storage address as an address for accessing the first data. Specification, page 64, lines 13-23.

VI. ISSUES

1. Claims 23-30 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 6,195,680 to *Goldszmidt* in view of U.S. Patent No. 5,829,033 to *Hagersten*.

VII. ARGUMENT

Claims 23-30 are not properly rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 6,195,680 to *Goldszmidt* in view of U.S. Patent No. 5,829,033 to *Hagersten*.

Claim 23 recites a method of operating a distributed processing system having a network coupling a multiplicity of Host distributed devices for processing workloads for the distributed processing system, a plurality of Client systems requesting processing of the workloads, and a Server system for selectively distributing the workloads from the plurality of Client systems for processing by the distributed processing system comprising five steps. In step 1, the Server system receives a request from one of the plurality of Client systems to use the distributed processing system to process a first workload. In step 2, the first workload is sent to a first Host distributed device selected from the multiplicity of Host distributed devices. In step 3, an index of one or more data addresses defining a location of first data required to process the first workload is sent to the first Host distributed device. In step 4, the first data from a first data address selected from the one or more data addresses in the index is accessed; and in step 5, the index is updated to include a storage address of storage coupled to the first Host distributed device as a location of the first data.

The Examiner states that *Goldszmidt* teaches the distributed processing system of Claim 23 and cites *Goldszmidt*, column 3, lines 12-21 and column 5, lines 55-63. The invention of *Goldszmidt* is directed to a client based system for the fault tolerant delivery of real-time data streams such as live audio or video. *Goldszmidt* does not

teach or suggest a method of a distributed processing system wherein a Server distributes workloads from Client systems to Host distributed devices for processing by the distributed processing system. The Client systems in *Goldszmidt* receive request data from a primary server, detects imbalances or failures in receiving the data and automatically switches to a secondary Server to minimize service disruption. See the Abstract of *Goldszmidt*. In column 3, lines 12-21, *Goldszmidt* describes an embodiment where switching between Servers is done in such a way that the switching is transparent. Nowhere in this recitation does *Goldszmidt* teach or suggest the distributed processing system of Claim 23.

In column 5, lines 55-63, *Goldszmidt* describes his protocols for switching to secondary Servers so that a Client system receives data with minimum distribution. Nowhere in this recitation does *Goldszmidt* teach or suggest the distributed processing system of Claim 23. The invention of *Goldszmidt* is directed to selectively switching between Servers systems to maintain a real-time data transmission and is not directed to the distributed processing system of Claim 23.

The Examiner states that *Goldszmidt* teaches step 1 of Claim 23 and cites column 3, lines 12-21 and column 4, lines 54-65. In step 1 of Claim 23 the Server system receives a request from one of the Client systems to use the distributed processing system to process a first workload. The Applicant has shown that in column 3, lines 12-21, *Goldszmidt* describes an embodiment where switching between Servers is done in such a way that the switching is transparent. *Goldszmidt* does not teach or suggest receiving a request from a Client system to process a workload on the distributed processing system. In column 4, lines 54-65, *Goldszmidt* describes that the control server that receives the requests from the Client systems for data could be a gateway through which the Client requests must and which includes a routing function to distribute the Client requests among servers in the cluster (of streaming Servers). The Servers of *Goldszmidt* are not receiving requests from Client systems to process workloads on the distributed processing system, rather *Goldszmidt* is describing how the Clients are requesting data from clusters of streaming Servers

configured to provide data (e.g., real-time audio and video) to minimize loss of content. Nowhere in this recitation does *Goldszmidt* teach or suggest step 1 of Claim 23.

The Examiner states that *Goldszmidt* teaches step 2 of Claim 23 and cites column 5, lines 33-47. Step 2 includes the step that the first workload is sent to a first Host distributed device selected from the multiplicity of Host distributed devices. *Goldszmidt* does not have a multiplicity of Host distributed devices coupled by a network to process workloads for the distributed processing system. *Goldszmidt* does not send workloads to be processed, *Goldszmidt* sends requests from a Client system for data provided by primary and secondary Servers in such a way to minimize real-time delivery of the data. In column 5, lines 33-47, *Goldszmidt* describes how the control server redirects Client requests to streaming Servers providing multimedia data. Nowhere in this recitation does *Goldszmidt* teach or suggest step 2 of Claim 23.

The Examiner admits that *Goldszmidt* does not specifically teach the use of a data address index as recited in step 3 of Claim 23. However the Examiner states that *Hagersten* teaches "a method of accessing the first data from a first data address selected from the one or more data addresses in the index" and cites column 24, lines 17-23, column 7, lines 1`8-29, and column 11, lines 14-21. The present invention the multiplicity of Host distributed devices coupled to process workloads for the distributed processing system may be widely separated. When a Client request that a workload be processed by the distributed processing system, the Server partitions the workload into workload units that are sent to the Host distributed devices for processing. The Host distributed devices may also require data to complete the task of processing the workload units. Initially the Server must provide the data by sending a data address index that defines where a Host distributed device can find the data necessary for its particular workload unit. In step 3 of Claim 23, an index of one or more data addresses defining a location of first data required to process the first workload is sent to the first Host distributed device. Claim 23 is not reciting just any data index, rather, the index of one or more data address defining the location of first

data required to process the first workload. The invention of *Hagersten* is directed to particular protocols for accessing shared data within a distributed memory system. The present invention is not concerned with managing memory devices, rather the present invention is concerned with providing a Host distributed device with an index defining the nearest location where the Host distributed device can find required data to process a particular workload unit. *Goldszmidt* does not teach or suggest the distributed processing system of Claim 23 nor does *Goldszmidt* teach or suggest steps 1-3. The Examiner did not show that *Hagersten* teaches or suggests the distributed processing system of Claim 23 nor steps 1-3. The Examiner states that *Hagersten* teaches accessing first data by using an address selected from one or more addresses in an index and cites column 7, lines 18-29, column 8, lines 4-20, and column 24, lines 17-23. In column 7, lines 18-29, *Hagersten* describes details of a typical address controller according to his discloser. In column 8, lines 4-20 *Hagersten* describes how addresses are broadcast to multiple memory devices coupled to an address bus. In column 24, lines 17-23, *Hagersten* describes the operations of his IGNORE and SHARE signals used when accessing distributed memory. The Applicant asserts that the Examiner is picking and choosing bits and pieces out of *Hagersten* and *Goldszmidt* in an attempt to find the elements of the steps in Claim 23. *Goldszmidt* is disclosing an invention directed to providing streaming data to a Client system by switching between cluster of Servers using a controller server. *Hagersten* is disclosing an invention directed to optimizing response in a system using distributed memory. Neither of these disclosures is directed to a method of managing the distributed processing system of Claim 23 using all the steps of the method in Claim 23. No one of ordinary skill in the art can arrive at the invention of Claim 23 by combining the teachings of *Hagersten* and *Goldszmidt*. There would no motivation to look to these two disclosure to try to arrive at the present invention. The Applicant can only discern bits of information about the way *Hagersten* addresses his distributed memory system and how *Goldszmidt* switches between Servers providing streaming data to minimize disruption in delivering the streaming data in real-time to a Client system. The Applicant asserts that the Examiner has not

shown where *Hagersten* and *Goldszmidt*, singly or in combination, teach or suggest the distributed processing system and any of steps 1-4 of Claim 23.

The Examiner states that it would have been obvious to combine the teachings of *Hagersten* and *Goldszmidt* to arrive at step 5 of Claims 23. In step 5 of Claim 23, the index is updated to include a storage address of storage coupled to the first Host distributed device as a location of the first data. The invention of Claim 23 starts with a distributed processing system comprising a multiplicity of Host distributed devices coupled through a network to process workloads for the distributed processing system. Claim 23 then recites steps that result in an index that contains addresses for locations storing data required to process workload units which is continually updated with new locations for the data whenever a new Host distributed device accesses the data to process a workload unit and updates the index with an address defining it as a location for the data. In this manner, Host distributed devices do not have to always access the Server to acquire data needed to process a workload unit if the data is available in its own storage or in storage of another Host distributed device that is in closer proximity. *Hagersten* and *Goldszmidt* do not teach or suggest the steps of Claim 23. No one or ordinary skill in the art cannot uses the teachings of *Hagersten* and *Goldszmidt*, singly or in combination, to arrive at the invention of step 5 of Claim 23 because they have a totally different functionality. Neither of these two references discuss distributed processing system that uses a multiplicity of Host distributed devices to process workload units partitioned from a workload to accomplish the processing of the workload. Therefore, the Applicant respectfully asserts that the rejection of Claim 23 under 35 U.S.C. § 103(a) as being unpatentable over *Goldszmidt* in view of *Hagersten* is traversed for the above reasons.

Claim 24 is dependent from Claim 23 and contains all the limitations of Claim 23. The Applicant has shown that *Hagersten* does not teach or suggest the invention of Claim 23. Claim 24 adds the limitation that the multiplicity of Host distributed devices coupled to the network to process workloads for the distributed processing system are so coupled in response to an incentive supplied by the Server system. The

Examiner states that *Hagersten* teaches Claim 24 and recites column 4, lines 40-50. The following is a direct quote of this recitation: “For example, a common reason for asserting the IGNORE signal is that a request cannot be serviced within the local subsystem, and thus should not be ordered at this time. In the absence of global reordering, responding to an asserted transaction requesting data could result in providing the requestor with an invalid version of the desired data.” The Applicant fails to see how this cited disclosure reads on Claim 24. It adds nothing relative to the invention of Claim 23 and it makes no mention of coupling a multiplicity of Host distributed devices to a network to process workloads for the distributed processing system and it makes no mention of the Host distributed devices being coupled in response to incentives. The Applicant asserts that the Examiner has failed to make a *prima facie* case of obviousness for rejecting Claim 24 using the recitation in column 24, lines 40-50 of *Hagersten*. Therefore, the Applicant respectfully asserts that the rejection of Claim 24 under 35 U.S.C. § 103(a) as being unpatentable over *Goldszmidt* in view of *Hagersten* is traversed for the above reasons and for the same reasons as Claim 23.

Claim 25 is dependent from Claim 24 and contains all the limitations of Claim 24. Claim 25 adds the limitation that the incentive defines an advantage for the multiplicity of Host devices to couple to the network. The Examiner states that *Hagersten* teaches Claim 25 and cites column 4, lines 40-50 listed above relative to Claim 24. The Applicant has shown that *Hagersten* does not teach or suggest the invention of Claim 23 or Claim 24. The Applicant asserts that the Examiner has failed to make a *prima facie* case of obviousness for rejecting Claim 24 using the recitation in column 24, lines 40-50 of *Hagersten* since the Examiner has failed to specifically point out where in this reference the limitation of Claim 25 is found. Therefore, the Applicant respectfully asserts that the rejection of Claim 25 under 35 U.S.C. § 103(a) as being unpatentable over *Goldszmidt* in view of *Hagersten* is traversed for the above reasons and for the same reasons as Claims 23 and 24.

Amended Claim 26 is dependent from Claim 23 and contains all the limitations of Claim 23. The Applicant has shown that *Goldszmidt* does not teach or suggest the invention of Claim 23. Claim 26 adds the limitation that the first Host distributed device is selected to process the first workload in response to capability values in a capability vector for the first Host distributed device in a capability database coupled to the Server system. The Examiner states that *Goldszmidt* teaches Claim 24 and recites column 5, lines 55-64. In column 5, lines 55-64, *Goldszmidt* describes the protocol of how the controller Server assigns and redirects a client receiving streaming data to one of the streaming servers in either of the two groups (server clusters). Nowhere in this recitation are capability values, capability vectors, or capability database mentioned. The Applicant asserts that the Examiner has failed to make a *prima facie* case of obviousness for rejecting Claim 26 using the recitation in column 5, lines 55-64 of *Goldszmidt* since the Examiner has failed to specifically point out where in this reference the limitation of Claim 26 is found. Therefore, the Applicant respectfully asserts that the rejection of Claim 26 under 35 U.S.C. § 103(a) as being unpatentable over *Goldszmidt* in view of *Hagersten* is traversed for the above reasons and for the same reasons as Claim 23.

Claim 27 is dependent from Claim 23 and contains all the limitations of Claim 23. The Applicant has shown that *Goldszmidt* does not teach or suggest the invention of Claim 23. Claim 27 adds the limitation that the first Host distributed device is selected to process the first workload in part because a data address for the first data required to process the first workload in the index corresponds to a storage address for accessing storage coupled to the first Host distributed device. The Examiner states that *Hagersten* teaches Claim 27 and recites column 7, lines 18-30. In this recitation *Hagersten* describes his typical address controller 180 in further detail. In particular *Hagersten* describes the protocol for processing requests issued to Address Bus 60 according to his invention. *Goldszmidt* is not disclosing how the first Host distributed device is selected to process the first workload as recited in Claim 23. The invention of *Goldszmidt* is directed to delivering streaming data in real-time to Client systems by switching among cluster of Server systems. The Applicant asserts that the

Examiner has failed to make a *prima facie* case of obviousness for rejecting Claim 27 using the recitation in column 7, lines 18-29 of *Goldszmidt* since the Examiner has failed to specifically point out where in this reference the limitation of Claim 27 is found. Therefore, the Applicant respectfully asserts that the rejection of Claim 27 under 35 U.S.C. § 103(a) as being unpatentable over *Goldszmidt* in view of *Hagersten* is traversed for the above reasons and for the same reasons as Claim 23.

Claim 28 is an independent claim directed to a distributed processing system comprising several elements. A multiplicity of Host distributed devices coupled to a network to process workloads for the distributed processing system. A Server system coupled to the network for distributing workloads to selected of the Host distributed devices. A database coupled to the Server system for storing capability vectors having capability values defining an ability each of the Host distributed devices has for process workloads for the distributed processing system. An index stored in the database having one or more storage addresses defining storage locations for accessing data required to process workloads for the distributed processing system. Circuitry for accessing first data required for a first workload by a first Host distributed device processing the first workload using an address of the first data stored in the index, wherein the first Host distributed device stores the first data at a first data address when processing the first workload and circuitry in the first Host distributed device for automatically updating the index to include the first data address as a location for the first data.

The Examiner states that *Goldszmidt* teaches the first element of Claim 28; . The multiplicity of Host distributed devices coupled to a network to process workloads for the distributed processing system citing column 3, lines 12-21 and column 5, lines 55-63. The Applicant has shown relative to Claim 23 that *Goldszmidt* does not teach or suggest a method of a distributed processing system wherein a Server distributes workloads from Client systems to Host distributed devices for processing by the distributed processing system. The Client systems in *Goldszmidt* receive request data from a primary server, detects imbalances or failures in receiving

the data and automatically switches to a secondary Server to minimize service disruption. See the Abstract of *Goldszmidt*. In column 3, lines 12-21, *Goldszmidt* describes an embodiment where switching between Servers is done in such a way that the switching is transparent. Nowhere in this recitation does *Goldszmidt* teach or suggest a multiplicity of Host distributed devices coupled to a network to process workloads for the distributed processing system as recited in Claim 28. In column 5, lines 55-63, *Goldszmidt* describes his protocols for switching to secondary Servers so that a Client system receives data with minimum distribution. Nowhere in this recitation does *Goldszmidt* teach or suggest the multiplicity of Host distributed devices coupled to a network to process workloads for the distributed processing system of Claim 28. The invention of *Goldszmidt* is directed to selectively switching between Servers systems to maintain a real-time data transmission and is not directed to the distributed processing system of Claim 28.

The Examiner states that *Goldszmidt* teaches the second element of Claim 28; the Server system coupled to the network for distributing workloads to selected of the Host distributed devices citing column 5, lines 55-64. The Applicant has shown relative to element 1 of Claim 28 that the disclosure in column 5, lines 55-64 is not directed to the invention of Claim 28. The Servers systems in *Goldszmidt* are used to maintain a real-time data transmission and are not directed to the distributed processing system of Claim 28.

The Examiner states that *Goldszmidt* teaches the third element of Claim 28; the database coupled to the Server system for storing capability vectors having capability values defining an ability each of the Host distributed devices has for process workloads for the distributed processing system citing column 5, lines 55-64 and column 17, lines 45-60. The Applicant has shown relative to element 1 of Claim 28 that the disclosure in column 5, lines 55-64 is not directed to the invention of Claim 28. The Servers systems in *Goldszmidt* are used to maintain a real-time data transmission and are not directed to the distributed processing system of Claim 28. Nowhere in the recitation of column 5, lines 55-64 does *Goldszmidt* teach or suggest

a capability database as recited in Claim 28. In column 17, lines 45-60 *Goldszmidt* describes his WriteBlockIO transaction and his WriteIO transaction as well as other signals *Goldszmidt* uses to store and retrieve data from his distributed memory system. *Goldszmidt* does not mention any database in this recitation let alone a database used to store capability vectors according to Claim 28.

The Examiner states that *Goldszmidt* does not specifically teach element 4; the use of an index stored in the database having one or more storage addresses defining storage locations for accessing data required to process workloads for the distributed processing system. The Examiner then recites column 7, lines 18-29, column 11, lines 14-21 and column 24, lines 17-23 for no stated reason and without stating to which reference he is referring.

The Examiner lists element 5 and 6; circuitry for accessing first data required for a first workload by a first Host distributed device processing the first workload using an address of the first data stored in the index, wherein the first Host distributed device stores the first data at a first data address when processing the first workload and circuitry in the first Host distributed device for automatically updating the index to include the first data address as a location for the first data. The Examiner then recites column 7, lines 18-29, column 24, lines 17-23 and column 8, lines 4-20 after element 5 and column 36, lines 36-40 and column 7, lines 18-30 again for no stated reason and without stating to which reference he is referring. Simply listing passages of unknown references does not meet the requirements for making a *prima facie* case of obviousness. The Applicant therefore asserts that the Examiner has failed to make a *prima facie* case of obviousness for rejecting Claim 28 since the Examiner has failed to specifically point out where in the references *Goldszmidt* and *Hagersten* all the elements of Claim 28 are found. Therefore, the Applicant respectfully asserts that the rejection of Claim 28 under 35 U.S.C. § 103(a) as being unpatentable over *Goldszmidt* in view of *Hagersten* is traversed for the above reasons.

Amended Claim 29 is an independent claim directed to a computer program product operating within a Server system managing a distributed processing system having a program of instructions for performing four steps. The Examiner states that *Goldszmidt* teaches a computer program product operating within the Server managing the distributed processing system of Claim 29 and recites column 3, lines 12-21 and column 5, lines 55-63. The Applicant has shown relative to Claim 23 and Claim 28 that *Goldszmidt* does not teach or suggest the distributing processing system comprising a Server coupled with a network to a multiplicity of Host distributed devices to process workloads for the distributed processing system. The references of *Goldszmidt* cited by the Examiner do not mention a computer program product of any kind and also do not disclose the distributed processing system of Claim 29.

The Examiner states that *Goldszmidt* teaches first program step that configures a database in storage coupled to the Server system for storing and accessing capability vectors having capability values defining an ability each of a multiplicity of Host distributed devices coupled to the Server for processing workloads for the distributed processing system citing column 5, lines 55-64 of *Goldszmidt*. The Applicant has shown relative to Claims 26 and 28 that *Goldszmidt* does not teach or suggest the database of Claim 29. The Applicant therefore asserts that the Examiner has failed to make a *prima facie* case of obviousness for rejecting this limitation of Claim 29 since the Examiner has failed to specifically point out where in the references *Goldszmidt* and *Hagersten* program step 1 of Claim 29 is found.

The Examiner states that *Goldszmidt* does not specifically teach program step 2 of Claim 29; configuring an index in the database for storing addresses defining storage locations of data required to process each workload the Server system submits to the distributed processing system for processing. The Examiner then recites column 7, lines 18-29, column 11, lines 14-21 and column 24, lines 17-23 for no stated reason and without stating to which reference he is referring.

The Examiner lists program step 3 and 4; sending storage addresses of first data required to process a first workload from the index to a first Host distribute device when the first Host distributed device is selected to process the first workload and updating the index with a storage address of the first data within storage coupled to the first Host distributed device when the first Host distributed device is selected to process the first workload. The Examiner then recites column 7, lines 18-29, column 24, lines 17-23 and column 8, lines 4-20 after program step 3 and column 36, lines 36-40 and column 7, lines 18-30 after program step 4 again for no stated reason and without stating to which reference he is referring. Simply listing passages of unknown references does not meet the requirements for making a *prima facie* case of obviousness. The Applicant therefore asserts that the Examiner has failed to make a *prima facie* case of obviousness for rejecting Claim 29 since the Examiner has failed to specifically point out where in the references *Goldszmidt* and *Hagersten* all the program steps of Claim 29 are found. Therefore, the Applicant respectfully asserts that the rejection of Claim 29 under 35 U.S.C. § 103(a) as being unpatentable over *Goldszmidt* in view of *Hagersten* is traversed for the above reasons and for the same reasons as Claims 23, 26, and 28.

Claim 30 is an independent claim to a software agent operating within each of a multiplicity of Host distributed devices coupled to a network, the network configured to enable the Server system to selectively couple the multiplicity of Host distributed devices to process workloads for a distributed processing system, wherein the software agent comprises a program of instructions for performing four program steps. The Examiner rejected Claim 30 for the same reasons as Claim 23 without specifically addressing the fact that Claim 30 is directed to a software agent operating in each of the multiplicity of Host distributed devices and Claim 23 is directed to a method of managing a distributed processing system. Likewise all of the steps of Claim 30 are not the same as Claim 23. The Applicant therefore respectfully asserts that the Examiner has failed to make a *prima facie* case of obviousness over *Goldszmidt* in view of *Hagersten* for failing to specifically point out where the software agent and each program step it performs is found in the cited references.


Therefore, the Applicant respectfully asserts that the rejection of Claim 30 under 35 U.S.C. § 103(a) as being unpatentable over *Goldszmidt* in view of *Hagersten* is traversed for the above reasons and for the same reasons as Claims 23.

Respectfully submitted,

WINSTEAD SECHREST & MINICK P.C.

Patent Agent and Attorney for Appellants

By:



Richard F. Frankeny

Reg. No. 47,573

Kelly K. Kordzik

Reg. No. 36,571

P.O. Box 50784
Dallas, Texas 75201
(512) 370-2851

APPENDIX

23. A method operating a distributed processing system having a network coupling a multiplicity of Host distributed devices for processing workloads for the distributed processing system, a plurality of Client systems requesting processing of the workloads, and a Server system for selectively distributing the workloads from the plurality of Client systems for processing by the distributed processing system comprising the steps of:

receiving a request by the Server system from one of the plurality of Client systems to use the distributed processing system to process a first workload;

sending the first workload to a first Host distributed device selected from the multiplicity of Host distributed devices;

sending to the first Host distributed device an index of one or more data addresses defining a location of first data required to process the first workload;

accessing the first data from a first data address selected from the one or more data addresses in the index; and

updating the index to include a storage address of storage coupled to the first Host distributed device as a location of the first data.

24. The method of claim 23, wherein the multiplicity of Host distributed devices are coupled to the network in response to an incentive supplied by the Server system.

25. The method of claim 24, wherein the incentive defines an advantage for the multiplicity of Host distributed devices to couple to the network.

26. The method of claim 23, wherein the first Host distributed device is selected to process the first workload in response to capability values of a capability vector for the first Host distributed device stored in a capability database coupled to the server system.

27. The method of claim 23, wherein the first Host distributed device is selected to process the first workload in part because a data address for the first data required to

process the first workload in the index corresponds to a storage address for accessing storage coupled to the first Host distributed device.

28. A distributed data processing system comprising:

a multiplicity of Host distributed devices coupled to a network such that the Host distributed devices process workloads for the distributed processing system;

a Server system coupled to the network for distributing workloads to selected of the multiplicity of Host distributed devices;

a database coupled to the Server system for storing capability vectors having capability values defining an ability of each of the multiplicity of Host processing devices has for processing workloads for the distributed processing system;

an index stored in the database having one or more storage addresses defining storage locations for accessing data required to process workloads for the distributed processing system;

circuitry for accessing first data required for a first workload by a first Host distributed device processing the first workload using an address of the first data stored in the index, wherein the first Host distributed device stores the first data at a first data address when processing the first workload; and

circuitry in the first Host distributed device for automatically updating the index in the database to include the first data address as a location for the first data.

29. A computer program product operating within a Server system managing a distributed processing system, wherein the Server system is coupled to a network, the network configured to enable the Server system to be coupled to a multiplicity of Host distributed devices for processing workloads for the distributed processing system, the program product comprising a program of instructions for performing the program steps of:

configuring a database in storage coupled to the Server system for storing and accessing capability vectors have capability values defining an ability each of the multiplicity of Host distributed devices has for processing workloads for the distributed processing system;

configuring an index in the database for storing addresses defining locations of data required to process each workload the Server system submits to the distributed processing system for processing;

sending storage addresses of first data required to process a first workload from the index to a first Host distributed device when the first Host distributed device is selected to process the first workload; and

updating the index with a storage address of the first data within storage coupled to the first Host distributed device when the first Host distributed device is selected to process the first workload.

30. A software agent operating within each of a multiplicity of Host distributed devices coupled to a network, the network configured to enable a Server system to selectively couple the multiplicity of Host distributed devices to perform workloads for a distributed processing system, the software agent comprising a program of instructions for performing the program steps of:

receiving in a first Host distributed device selected from the multiplicity of Host distributed devices a first workload, an application program to process the first workload, and an index of storage addresses defining one or more locations for accessing first data required to process the first workload;

accessing the first data from one of the storage addresses in the index;

storing the first data at a first storage address in storage coupled to the first Host distributed device when the first Host distributed device is selected to process the first workload; and

updating the index by adding the first storage address as an address for accessing the first data.